

CLAIMS

1. An apparatus comprising a data set stored on machine-readable media, wherein:

the data set is divided into multiple subsets;

a spread memory layout is implemented on the machine-readable media, the spread memory layout defining multiple pages in memory, with a subset of data from the data set being mapped to one or more predetermined portions of each page, the portions being less than the capacity of each page; and

each page is sized to map to quick access memory of a processor, such that image data when fetched from the machine-readable media are mapped into one or more predetermined portions of quick access memory.

2. An apparatus according to claim 1, wherein the data set is primarily volumetric data, and wherein each subset is selected to be a cubic region of volumetric data, with all cubic regions being of like-size.

3. An apparatus according to claim 2, wherein each cubic region includes an 8x8x8 cube of voxels.

4. An apparatus according to claim 1, adapted for use where the quick access memory is an on-chip processor cache, said apparatus further being characterized in that the size of each page having any data from the data set corresponds to size of the on-chip processor cache.

5. An apparatus according to claim 1, adapted for use where the quick access memory is a processor cache, said apparatus further being characterized in that the size of each page having any data from the data set is selected to map to at least a portion of the processor cache.

6. An apparatus according to claim 4, wherein said apparatus is adapted for use with a n -way associative cache, said apparatus being further characterized in that size of each page having image data equals the size of the processor cache divided by n .

7. An apparatus according to claim 4, further comprising sizing each page such that a subset of data from the data set stored in that page occupies less than one-half of page size, and such that data when fetched from machine-readable media can only be loaded into predetermined lines of processor cache, no more than one-half of processor cache capacity.

8. An apparatus according to claim 7, wherein each page is sized such that a subset of data from the data set stored in that page occupies no more than approximately one-quarter of page size, and wherein each subset occupies one or more predetermined, consistent locations within a page, such that all subsets of data from the data set are mapped to one or more specific locations in processor cache each time a page is retrieved from the machine-readable media.

9. An apparatus according to claim 1, wherein each subset occupies a specific contiguous location within each page that is at least one kilobyte in size.

10. An apparatus according to claim 1, wherein each page is at least four kilobytes in size.

11. An apparatus according to claim 1, adapted for use where the data set substantially consists of image data, further comprising processing parameters selected from the group of buffers, tables, a tile buffer, shading tables, and rendering parameters, the processing parameters being stored in a page of memory at a location other than the one or more predetermined portions, such as to inhibit overwrite of processing parameters by image data.

12. An apparatus according to claim 1, wherein the data for each subset is stored on the machine-readable media in a dimensionally-interleaved manner.

13. An apparatus according to claim 12, wherein the image data is volumetric data and for each subset is stored on the machine-readable media in a cubically-interleaved manner.

14. An apparatus according to claim 12, adapted for use where the data set represents three or more dimensions, wherein the data for each subset is stored on the machine-readable media in a dimensionally-interleaved manner where data for at least three-dimensions are interleaved.

15. An apparatus for processing a data set, comprising:

means for mapping data from the data set as part of a page to one or more predetermined portions of a cache; and

means for mapping processing parameters as part of a page to predetermined portions of the cache, such that processing parameters occupy different portions of cache than data from the data set;

wherein the addressing of data from the data set and processing parameters is thereby structured such that processing parameters do not overwrite data from the data set when loaded into the cache and such that data from the data set does not overwrite processing parameters when loaded into the same cache.

16. An apparatus according to claim 15, further comprising:

means for dimensionally-interleaving data from the data set as discrete, contiguous subsets of data, such that each subset represents a multidimensional space.

17. An apparatus according to claim 15, adapted for processing volumetric data, said apparatus further comprising:

means for cubically-interleaving volumetric data representing three-dimensions using oct-tree addressing.

18. A method of processing a data set, comprising:

storing the data set in a data set memory in a spread memory layout, with at least a predetermined gap in memory address between each one of multiple subsets of the data;

storing processing parameters selected in one or more predetermined parts of cache;

retrieving portions of the image data and mapping retrieved image data into the cache;

wherein an addressing scheme implemented by the spread memory layout and wherein the mapping into cache cause retrieved data from the data set to be loaded into the cache at locations other than the predetermined parts of cache.

19. A method according to claim 18, wherein the data set substantially consists of volumetric data, further comprising selecting each subset to be a cubic region of volumetric data, with all cubic regions being of like-size.

20. A method according to claim 19, further comprising selecting each cubic region to be an 8x8x8 cube of voxels.

21. A method according to claim 18, adapted for use where the quick access memory is an on-chip processor cache, said method further comprising sizing a gap in memory address in correspondence to a predetermined mapping to on-chip processor cache, such that data from the data set is thereby mapped to at least one predetermined space within the on-chip processor cache.

22. A method according to claim 18, adapted for use where the quick access memory is a processor cache, said method further comprising sizing a gap in memory address in correspondence to a predetermined mapping to the processor cache, such that data is thereby mapped to at least one predetermined space within the processor cache.

23. A method according to claim 22, wherein said method is adapted for use with a n -way associative cache, said method further comprising sizing a gap in memory address such that memory subsets are repeated within memory at an interval equal to the size of the processor cache divided by n .

24. A method according to claim 22, further comprising sizing the gap each in memory address such that each subset of data from the data set when loaded into cache occupies less than one-half of a direct or n -way associated cache space, and such that data from the data set when fetched from machine-readable media can only be loaded into predetermined lines of the cache space, the predetermined lines representing no more than one-half of processor cache capacity.

25. A method according to claim 24, further comprising sizing a gap in memory address such that a subset of data from the data set occupies no more than approximately one-quarter of mapped cache space.

26. A method according to claim 18, further comprising implementing the spread memory layout results in a mapping of data from the data set to a cache location that is at least one kilobyte in size.

27. A method according to claim 18, further comprising storing each subset of data from the data set in a dimensionally-interleaved manner.

28. A method according to claim 27, and further adapted for the processing of volumetric data, further comprising storing each of multiple subsets of volumetric data in a cubically-interleaved manner.

29. An apparatus, comprising:

a processor;

a quick access memory adapted for use by the processor in processing a data set;

an offline mechanism for storing the data set in a main memory using a spread memory layout, wherein the spread memory layout is chosen such that

the data set is divided into subsets of data,

the subsets are each stored in a page of memory of predetermined size, the page size chosen to map to at least a predetermined portion of the quick access memory, and

the subsets are restricted to occupy only a selected portion of each page; and

instructions stored on machine readable media that cause said processor to load processing parameters used for processing the data set into the predetermined portions of quick access memory;

wherein the selected portion of each page is chosen such that the subsets of data will only occupy specific memory locations within mapped portion of quick access memory, and wherein the instructions cause the loading of processing parameters into the mapped portion of quick access memory at locations other than the specific memory locations, such that the data set is inhibited from overwriting processing parameters when data from the data set is retrieved from main memory and loaded in mapped fashion into the quick access memory.

30. An apparatus according to claim 29, wherein the data set represents multiple dimensions and wherein the offline mechanism further stores data from the data set in a dimensionally-interleaved manner, such that each page in memory includes dimensionally-interleaved data.

31. An apparatus according to claim 30, wherein the data set substantially consists of volumetric data and wherein the offline mechanism further stores data from the data set in a cubically-interleaved fashion.

32. An apparatus according to claim 30, wherein the data set substantially consists of image data and wherein the offline mechanism further stores image data from the data set in a dimensionally-interleaved manner, such that each page in memory includes data that is dimensionally-interleaved in two or more dimensions.

33. An apparatus according to claim 29, wherein the quick access memory is a n -way processor cache and wherein page size is selected to map the subsets to a predetermined portion of any of n sections of processor cache.

34. An apparatus according to claim 33, wherein the quick access memory is an on-chip processor cache.

35. An apparatus according to claim 29, wherein the data set substantially consists of volumetric data, said apparatus further comprising instructions stored on machine-readable media that when executed cause said processor to perform raycasting using the data set.

36. An apparatus according to claim 29, wherein the processing parameters include parameters selected from the group of a tile buffer, a pixel buffer, shading coefficients, lighting coefficients, texture information and light source information.